DSP Block Design Proposal for PICo: Team NOR

Matthew Ridder, Joseph Fackler, Doug Etts
ECE 3663 – Spring 2014
University of Virginia
<mjr3vk, jmf6ew, dee8ak>@virginia.edu

ABSTRACT
In this paper, Team NOR lies out and explains the rationale behind their design for the specified DSP block.

1. INTRODUCTION
The DSP block requested by PICo requires the capacity to compute eight separate functions: NOP, PASS A, ADD, SUB, SHIFT, AND, OR, and one function of the designer's choosing. The block must perform these operations on two separate 16 bit inputs, A<15:0> and B<15:0>. Three control bits CON<2:0> determine which function the block should output. The expected output, OUT<15:0>, is determined by the function that was indicated by the control inputs. The DSP block is synchronized by a CLK input through the use of three registers: one at A<15:0>, one at B<15:0>, and one at ALUout<15:0>.

After attaining functionality, Team NOR focused on minimizing the PICo design metric. This metric is determined by the product of the energy per cycle, delay, and area as determined by transistor width. The design decisions made to optimize this metric will be touched on throughout the paper.

2. DESIGN DESCRIPTION
A top level view of the DSP block will be discussed in this section.

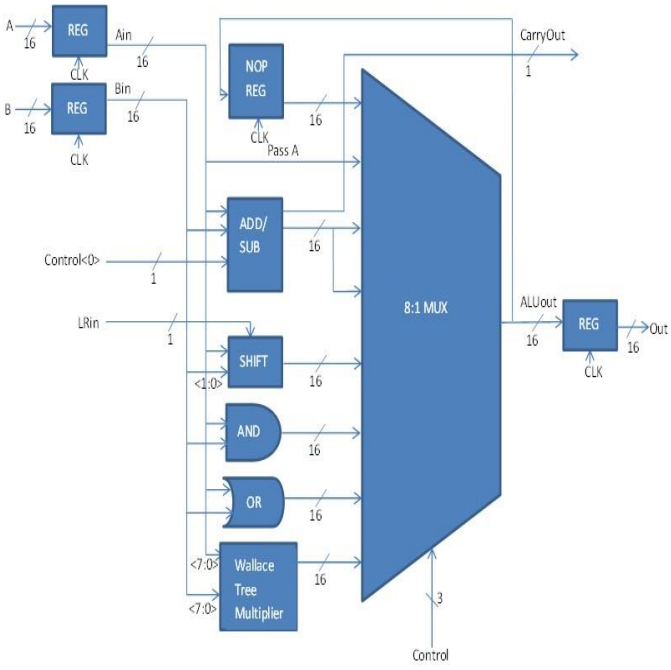

Figure 1. Top Level Block Diagram.

2.1 NOP: No Operation
The NOP function simply preserves the previous state of the output OUT<15:0> by returning the output of the ALU from the previous operation through a register. The NOP function is triggered by control CON<2:0> = 000b.

2.2 PASS A: Out = A
The PASS A function is composed of 16 parallel transmission gates which each output a bit of A<15:0>. The output OUT<15:0> is equal to input A<15:0>. PASS A is triggered by CON<2:0> = 001b.

2.3 ADD/SUB: Out = A+B/Out=A-B
The ADD and SUB functions are handled by a single mirror adder/subtractor block. The ADD function, triggered by CON<2:0> = 010b, outputs OUT<15:0> = A<15:0> + B<15:0>, with COUT<0> as a carry flag. The SUB function, triggered by CON<2:0> = 011b, outputs OUT<15:0> = A<15:0> - B<15:0>, with COUT<0> serving as a sign flag. In this case, when COUT<0> = 0, the output of the block is positive and when COUT<0> = 1, the output of the block is negative and the 2's complement must be taken to find the magnitude. Team NOR's combination of the ADD and SUB components in the DSP block into a single ADD/SUB block is made possible by the fact that addition and subtraction is very similar at the transistor level. The only additional gates needed to combine the components are XOR gates controlled by the "carry-in" bit of the adder in order to switch between the ADD and SUB functions. In Team NOR's ADD/SUB block, the "carry-in" of the adder is CON<0>. Each bit of B<15:0> is put into an XOR gate against the CON<0> bit, so when CON<0> = 1, triggering subtraction, B<15:0> undergoes bitwise inversion. This, combined with the carry-in, allows the SUB function to operate in two's complement. Due to this XOR-gate, COUT<0> serves as a sign flag for the output OUT<15:0> of the SUB function. This does not interfere with the operation of the ADD block, since CON<0> = 0 when the ADD function is at work. The output of the block attaches to both the 010 and 011 inputs of the ALU mux.

2.4 SHIFT: Out = A<<B
The SHIFT function uses a barrel shifter structure. Triggered by CON<2:0> = 100b, the output OUT<15:0> is a bit-shifted form of input A<15:0>. Input B<1:0> sets the magnitude of the shift, up to 4 bits, while an extra input, LRIN<0> sets the direction left or right. LRIN<0> = 0 indicates a right shift while LRIN<0> = 1 indicates a left shift.

2.5 AND: Out = A & B
The AND function is handled by 16 parallel 2-input AND gates, and the output OUT<15:0> is a bitwise AND of the respective bits in inputs A<15:0> and B<15:0>. The AND function is triggered by CON<2:0> = 101b.

2.6 OR: Out = A | B

The OR function is handled by 16 parallel 2-input OR gates, and the output OUT<15:0> is a bitwise OR of the respective bits in inputs A<15:0> and B<15:0>. The OR function is triggered by CON<2:0> = 110b.

2.7 Arbitrary: Out = A<7:0> * B<7:0>

The arbitrary function that Team NOR chose to implement is eight bit by eight bit multiplication with the two eight bit inputs taken from inputs A<15:0> and B<15:0>. The multiplication is handled by an eight bit by eight bit Wallace Tree multiplier structure. The multiplier is triggered by CON<2:0> = 111b.

3. INNOVATION

The methodology used to optimize the metric and the justification of design decisions will be discussed in this section.

3.1 Key Trade-offs

This subsection will discuss the key trade-offs made in some of the components within the DSP block.

3.1.1 PASS A

Transmission gates are used instead of pass gates for this function. This allows the full range of voltage values from the rails to be passed through at the cost of an increase in area. This tradeoff is worth it because it minimizes the risk of functionality failure later on in the DSP block by passing strong '0's and '1's.

3.1.2 SHIFT

Team NOR utilizes a barrel shifter structure for the implementation of the SHIFT function. With this implementation, the signal that is being shifted has to pass through at most one transmission gate. While this is a significant benefit for optimizing delay, there is a slight increase in area cost resulting from the decoder that is required to turn on and off specific gates to pass through the correct signals to the output. This trade-off is beneficial to the DSP block because the increase in area in the shifter is nominal compared to the overall size of the block. The delay here is much more important to the overall performance. Also, transmission gates are used instead of pass gates to allow the full range of voltage values to be passed through. Again, this results in an area cost, but it is worth it in order to minimize the risk of functionality failure.

Team NOR's SHIFT function has the added capability of being able to shift left in addition to being able to shift right. This capability was added because shifters are commonly used in processors to multiply and divide by powers of 2. There is a large area cost to make this possible, but because this functionality is so common, the increase in area is worth it.

3.1.3 ADD/SUB

The ADD/SUB block is an integral piece of Team NOR's design. It contains the critical path of the entire DSP block and is also the bedrock of the Wallace Tree multiplier. Because of this, the delay became the main focus of metric optimization for this block. This metric is optimized through sizing the transistors within the block.

Team NOR's ADD/SUB block uses minimum size transistors for the sum stages in the mirror blocks and three times the minimum size transistors for the carry stages. Three times the minimum size was chosen in order to maintain optimal fanout within both parts of the mirror block. Before choosing this sizing scheme, Team NOR had tried sizing the blocks with logical effort and found that while this did improve the worst case delay of the ADD function, the SUB function delay was nearly 100% worse than it was without logical effort. This behavior is demonstrated in Figures 2 and 3. The signals shown indicate when the last sum bit settles after the worst case delay. The switching input, A<15> (not shown), triggers the worst case delay at around .5 ns. The red signals are generated by an ADD/SUB block using the minimum/fanout-optimal sizing scheme, and the blue signals are generated by an ADD/SUB block using the logical effort sizing scheme. Similar results were found when a Manchester Carry-Chain structure was implemented. Since the SUB function contains the critical path, the decision was made to resort back to minimum/fanout-optimal sizing at the cost of an increase in the worst case delay of the ADD function.

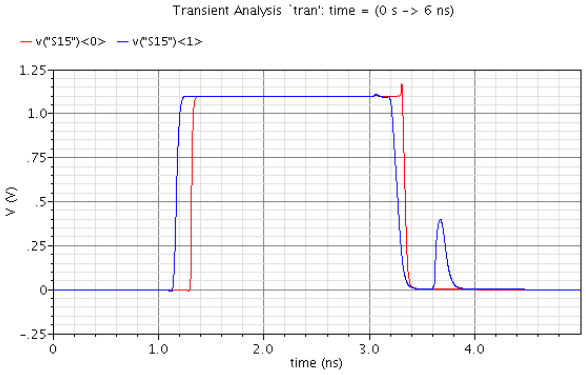

Figure 2. Logical Effort Sizing comparison for the ADD function.

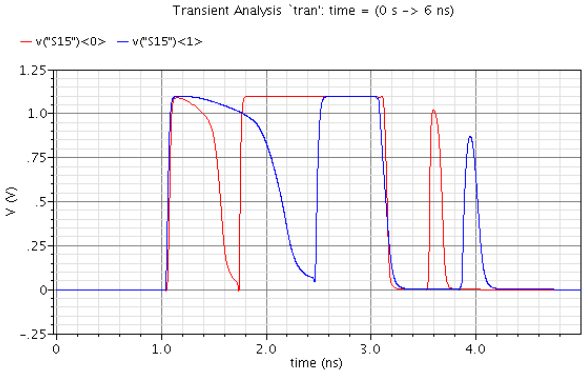

Figure 3. Logical Effort Sizing comparison for the SUB function.

Team NOR's combination of the ADD and SUB components in the DSP block conserves a large amount of area. Although this combination does require a few XOR gates, the area cost of those gates doesn't even come close to the area cost of creating an entirely separate SUB block.

Team NOR's ADD/SUB block is built as a string of mirror adders with necessary inverters built into the block to allow for each mirror adder's internal inversions. A mirror adder structure was chosen in order to avoid the large PMOS stacks and redundant

inverters that are inherent in static CMOS adder blocks. This structure also reduces the area of the block.

3.2 Sizing Methodology

The critical path that triggers the worst case delay of the metric is through the ADD/SUB block. Because of this, Team NOR made sure to design that block to optimize the delay. For the rest of the DSP block, the area and active power metrics were taken more into consideration when sizing each individual component because these components are not along the critical path. To optimize these metrics, minimum sizing to match a characteristic inverter is used. This minimizes area and power consumption for every component outside of the ADD/SUB block without harming the worst case delay of the DSP block.

3.3 Arbitrary Function

Multiplication was chosen as the arbitrary function due to its importance and prevalence in digital signal processing applications. Team NOR handles multiplication through the use of an eight bit by eight bit Wallace Tree multiplier. Since multiplication is essentially a bunch of additions of the partial products, a Wallace Tree structure was chosen because the adders can be arranged in a treelike fashion which can both reduce the critical path and the number of adders needed. This lowers both the delay of the multiplier and the total area of the multiplier. Once the tree compresses the partial products, a fast adder determines the final output. This output is 16 bits and the carry out bit is always 0. Team NOR uses a copy of their mirror adder that is used for the ADD/SUB function of their ALU with the carry-in tied to 0 for the fast adder in order to optimize the speed of the final addition.

4. RESULTS

Here is the metric for the DSP block: **7.92E-22 s*m*J**

The bolded values in each table are used in the calculation of this metric.

Table 1. Delay breakdown for each component

Component/Operation	Delay (s)
Register	4.56E-11
Pass A	3.49E-12
Add	1.07E-10
Sub	3.67E-10
Shift	4.69E-11
And	2.99E-11
Or	2.62E-11
ALU Critical Path	**4.64E-10**

Table 2. Average Power/Energy breakdown for each component

Component	Power (W)	Energy (J)
A Register	0.000682	1.0912E-09
B Register	0.000006	9.6E-12
NOP Register	0.000005	8E-12
Pass A	0.000004	6.4E-12
Add/Sub	0.000081	1.296E-10
Shift	0.000107	1.712E-10
And	0.000012	1.92E-11
Or	0.000011	1.76E-11
Mux	0.00003	4.8E-11
Out Register	0.000005	8E-12
Total	0.000943	**1.5088E-09**

Table 3. Area breakdown for each component

Component	Area (m)
Register	9.936E-5
Pass A	8.64E-6
Add/Sub	3.0105E-4
Shift	9.153E-5
And	1.584E-5
Or	1.872E-5
Mux	2.9808E-4
Total	**.0011313**

Table 4. Results for the Wallace Tree Multiplier

Metric	Value
Delay	1.02E-10 s
Area	5.4396E-4 m
Average Power	0.009726 W

5. CONCLUSION

PICo should select Team NOR's design for a few reasons. The first is that their design provides a wide range of functions that are commonly used in digital signal processing applications. The most notable of these are the left and right shift capabilities in the shifter, and the multiplier. Another reason is that Team NOR has ensured that the functionality of each component will always perform as expected since it uses transmission gates to pass a full range of voltage values. Team NOR has also designed their DSP block in a way that minimizes power, area, and delay. Along the critical path, which is in the ADD/SUB block, the delay metric is focused on and optimized. In that block, a sizing scheme is used

that provides a good compromise between the worst case delays for both the ADD and SUB functions. Also, because the ADD and SUB components are combined into a single block, a ton of area is saved. Outside of the critical path, the power and area metrics are optimized by using minimum size transistors. This significantly reduces power and area without harming the worst case delay of the DSP block. This focused optimization in different parts of the DSP block gives Team NOR's design a good balance of delay, area, and power. No one metric is far worse than any of the others. All of these reasons make Team NOR's DSP block a great design to choose.